Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
		(join same "first table" same "second table" and @ad<"20010523") and (ind\$3 or key\$) and distribut\$3 and resqust\$3 and stor\$3 and access\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 17:04
		((join same "first table" same "second table" and @ad<"20010523") and stor\$3 same row\$1) and column\$ with ("first table" and "second table")	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/24 17:18
L1	31	("first storage" or "storage one") and ("second storage" or "storage two") and table\$1 and row\$1 and join! and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/16 12:31
L2	93	("first storage" or "storage one" or "first database" or "database one") and ("second storage" or "storage two" or "second database" or "database two") and table\$1 and row\$1 and join! and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/16 14:35
L3	398	distribut\$3 with row\$1 and join! and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/16 14:35
L4	7	3 and ("first storage" or "storage one" or "first database" or "database one") and ("second storage" or "storage two" or "second database" or "database two")	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/16 14:36
S1	0	stor\$3 and row\$1 and "first table" with first near10 ("access module" or AMP) and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/24 17:11
S2	5	stor\$3 and row\$1 and "second table" and second near10 ("access module" or AMP) and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/24 17:12
S3	73	join same "first table" same "second table" and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/25 10:28
S4	47	(join same "first table" same "second table" and @ad<"20010523") and stor\$3 same row\$1	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/24 17:16

S5	18	((join same "first table" same "second table" and @ad<"20010523") and stor\$3 same row\$1) and column\$1 with ("first table" and "second table")	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 15:54
S6	0	(join same "first table" same "second table" and @ad<"20010523") and (ind\$3 or key\$1) and distribut\$3 and resqust\$3 and stor\$3 and access\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/24 18:02
S7	0	(join same "first table" same "second table" and @ad<"20010523") and (ind\$3 or key\$1) and resqust\$3 and stor\$3 and access\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/24 18:03
S8	0	(join same "first table" same "second table" and @ad<"20010523") and (ind\$3 or key\$1) and resqust\$3 and stor\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/24 18:04
S9	54	(join same "first table" same "second table" and @ad<"20010523") and (ind\$3 or key\$1) and stor\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/24 18:04
S10	1	("access module" or AMP) and join same "first table" same "second table" and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 12:06
S11	79	("access module" or AMP) and join same table\$1 and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/25 10:27
S12	102	join same "first table" same "second table" and stor\$3 and row\$1 and column\$1 "spool table" and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/25 10:39
S13	0	join same "first table" same "second table" same stor\$3 same row\$1 same column\$1 and "spool table" and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/25 10:37
S14	39	join same "first table" same "second table" same stor\$3 same row\$1 same column\$1 "spool table" and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/25 10:33
S15	0	join same "first table" same "second table" and stor\$3 and row\$1 and column\$1 and "spool table" and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2003/06/25 10:39

S16	0	join same "first table" same "second table" and "spool table" and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO;	OR	ON	2003/06/25 10:40
		and wad 20010323	DERWENT			
S22	1	"5379419".PN.	USPAT; USOCR	OR	ON .	2005/12/15 10:33
S23	1	"5437032".PN.	USPAT; USOCR	OR	ON	2005/12/15 10:33
S24	1	"5530939".PN.	USPAT; USOCR	OR	ON	2005/12/15 10:34
S25	257	teradata	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 12:05
S26	1	S24 and distribut\$3 and storage\$1 near module\$1 and table\$1 and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 12:23
S27	1	S24 and distribut\$3 and storage\$1 and table\$1 and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 12:10
S28	2	"6944633".pn.	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 13:02
S29	1	"5241648".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:27
S30	1	"5551031".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:27
S31	1	"5557791".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:29
S32	1	"5742806".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:29
S33	1	"5983215".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:32
S34	1	"6226639".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:34
S35	1	"6253197".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:37
S36	1	"6289334".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:37
S37	1	"6453314".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:38
S38	1	"6505189".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:38

S39	1	"6665684".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:40
S40	1	"5241648".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:42
S41	1	"5440730".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:42
S 4 2	1	"5666525".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:43
S43	1	"5241648".PN.	USPAT; USOCR	OR	ON .	2005/12/15 12:43
S44	1	"5345585".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:47
S45	1	"5537589".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:47
S46	1	"5557791".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:55
S47	1	"5710915".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:55
548	1	"5870747".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:56
S49	1	"5987453".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:56
S50	1	"6289334".PN.	USPAT; USOCR	OR	ON	2005/12/15 12:59
S51	0	EP1220098	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 13:02
S52	11	"1220098"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 13:02
S53	5	((join same "first table" same "second table" and @ad<"20010523") and stor\$3 same row\$1) and column\$1 with ("first table" and "second table") and "first table" and "second table" and module\$1	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 17:07
S54	45	stor\$3 same row\$1 same "first table" same storage\$1 and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 15:56
S55	32	stor\$3 same row\$1 same "second table" same storage\$1 and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 15:56

S56	27	S54 and S55	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 16:06
S57	14	S54 and S55 and join!	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 16:06
S59	10469	"first storage" same "second storage" and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR:	ON	2005/12/16 12:27
S60	1	S59 and (stor\$3 same row\$1) and column\$1 with ("first table" and "second table") and "first table" and "second table" and module\$1	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 17:38
S61	1	"4078260".PN.	USPAT; USOCR	OR	ON	2005/12/15 17:11
S62	0	S59 and (stor\$3 same column\$1) and row\$1 with ("first table" and "second table") and "first table" and "second table" and module\$1	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 17:39
S63	1026	S59 and (stor\$3 same column\$1)	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/16 09:10
S64	1	S63 and "first table" and "second table" and module\$1	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/15 17:40
S65	6	S63 and "first table" and "second table"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/16 09:11
S66	10469	"first storage" same "second storage" and @ad<"20010523"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/16 09:11
S67	976	S66 and (stor\$3 same row\$1)	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/16 09:11
S68	49	S66 and "first table" and "second table"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/16 09:13
S69	8	S67 and "first table" and "second table"	US-PGPUB; USPAT; EPO; JPO; DERWENT	OR	ON	2005/12/16 09:16

S70	8	S67 and "first table" and "second	US-PGPUB;	OR	ON	2005/12/16 09:17
		table" and row\$1 and stor\$3 and	USPAT;			
		column\$1 and "first table" and	EPO; JPO;			
		"second table"	DERWENT			



Subscribe (Full Service) Register (Limited Service, Free) Login

Search:

The ACM Digital Library O The Guide

distribute rows in a first storage and second storage



THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used distribute rows in a first storage and second storage

Found 103,167 of 167,655

Sort results by

Display

results

relevance expanded form

Save results to a Binder 3 Search Tips ☐ Open results in a new

Try an Advanced Search Try this search in The ACM Guide

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10

Relevance scale 🔲 📟 📟

Best 200 shown

Evolving RPC for active storage

Muthian Sivathanu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

window

October 2002 ACM SIGPLAN Notices, ACM SIGOPS Operating Systems Review, ACM SIGARCH Computer Architecture News, Proceedings of the 10th international conference on Architectural support for programming languages and operating systems ASPLOS-X, Volume 37, 36, 30 Issue 10, 5, 5

Publisher: ACM Press

Full text available: pdf(1.56 MB)

Additional Information: full citation, abstract, references

We introduce Scriptable RPC (SRPC), an RPC-based framework that enables distributed system services to take advantage of active components. Technology trends point to a world where each component in a system (whether disk, network interface, or memory) has substantial computational capabilities; however, traditional methods of building distributed services are not designed to take advantage of these new architectures, mandating wholesale change of the software base to exploit more powerful hardw ...

2 LH*RS: a high-availability scalable distributed data structure using Reed Solomon



<u>Codes</u>

Witold Litwin, Thomas Schwarz

May 2000 ACM SIGMOD Record, Proceedings of the 2000 ACM SIGMOD international conference on Management of data SIGMOD '00, Volume 29 Issue 2

Publisher: ACM Press

Full text available: pdf(155.52 KB)

Additional Information: full citation, abstract, references, citings, index terms

LH*RS is a new high-availability Scalable Distributed Data Structure (SDDS). The data storage scheme and the search performance of LH*RS are basically these of LH*. LH*RS manages in addition the parity information to tolerate the unavailability of k ⪈ 1 server sites. The value of k scales with the file, to prevent the reliability decline. The parity calculus uses the Reed -Solomon Codes. The storage and access performance over ...

Keywords: Reed-Solomon Codes, SDDS, high-availability, scalable

3 <u>LH*_{RS}---a highly-available scalable distributed data structure</u> Witold Litwin, Rim Moussa, Thomas Schwarz





September 2005 ACM Transactions on Database Systems (TODS), Volume 30 Issue 3

Publisher: ACM Press

Full text available: pdf(774.32 KB) Additional Information: full citation, abstract, references, index terms

LH* RS is a high-availability scalable distributed data structure (SDDS). An LH* RS file is hash partitioned over the distributed RAM of a multicomputer, for example, a network of PCs, and supports the unavailability of any k ≥ 1 of its server nodes. The value of k transparently grows with the file to offset the reliability decline. Only the number of the storage nodes potentially limits the file growth. The high-availability management uses a novel ...

Keywords: P2P, Scalable distributed data structure, grid computing, high-availability, linear hashing, physical database design

A provably efficient algorithm for dynamic storage allocation



E G Coffman, F T Leighton

November 1986 Proceedings of the eighteenth annual ACM symposium on Theory of computing

Publisher: ACM Press

Full text available: Told pdf(1.41 MB)

Additional Information: full citation, references, citings, index terms

5 System design methodologies and experiences: Low power storage cycle budget distribution tool support for hierarchical graphs



Erik Brockmeyer, Arnout Vandecappelle, Sven Wuytack, Francky Catthoor September 2000 Proceedings of the 13th international symposium on System synthesis

Publisher: IEEE Computer Society

Full text available: pdf(110.75 KB) Additional Information: full citation, abstract, references, citings

In data dominated applications, like multi-media and telecom applications, data storage and transfers are the most important factors in terms of energy consumption, area and system performance. Several steps which optimize these costs are present in our systematic Data Transfer and Storage Exploration methodology. In the important step discussed in this paper, the cycle budget available for background storage transfers is globally distributed over the application's memory accesses that are typic ...

6 Distributed operating systems

Andrew S. Tanenbaum, Robbert Van Renesse

December 1985 ACM Computing Surveys (CSUR), Volume 17 Issue 4

Publisher: ACM Press

Full text available: pdf(5.49 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Distributed operating systems have many aspects in common with centralized ones, but they also differ in certain ways. This paper is intended as an introduction to distributed operating systems, and especially to current university research about them. After a discussion of what constitutes a distributed operating system and how it is distinguished from a computer network, various key design issues are discussed. Then several examples of current research projects are examined in some detail ...

On the storage and retrieval of continuous media data

Banu Özden, Rajeev Rastogi, Avi Silberschatz

November 1994 Proceedings of the third international conference on Information and knowledge management

Publisher: ACM Press

Full text available: pdf(857.76 KB) Additional Information: full citation, abstract, references, index terms

Continuous media applications, which require a guaranteed transfer rate of the data, are becoming an integral part of daily computational life. However, conventional file systems do not provide rate guarantees, and are therefore not suitable for the storage and retrieval of continuous media data (e.g., audio, video). To meet the demands of these new applications, continuous media file systems, which provide rate guarantees by managing critical storage resources such as memo ...

Status report of the graphic standards planning committee



Computer Graphics staff

August 1979 ACM SIGGRAPH Computer Graphics, Volume 13 Issue 3

Publisher: ACM Press

Full text available: pdf(15.01 MB) Additional Information: full citation, references, citings

9 Analysis of interleaved storage via a constant-service queuing system with Markov-





chain-driven input

Micha Hofri

June 1984 Journal of the ACM (JACM), Volume 31 Issue 3

Publisher: ACM Press

Full text available: mpdf(1.02 MB)

Additional Information: full citation, abstract, references, index terms, review

A popular means of increasing the effective rate of main storage accesses in a large computer is a multiplicity of memory modules accessible in parallel. Although such an organization usually achieves a net gain in access rate, it also creates new modes of congestion at the storage controller. This paper analyzes the variables that describe such a congestion: queue lengths and delays. A controller that maintains separate register sets to accommodate the request queue of each module is ...

Keywords: Congestion, batch arrival, constant service, queuing system, waiting times

10 Semi-linear and bi-base storage schemes classes: general overview and case study



J. Jorda, A. Mzoughi, D. Litaize

July 1995 Proceedings of the 9th international conference on Supercomputing

Publisher: ACM Press

Full text available: pdf(839.85 KB) Additional Information: full citation, references, index terms

11 The design, implementation, and evaluation of a symmetric banded linear solver for



distributed-memory parallel computers

Anshul Gupta, Fred G. Gustavson, Mahesh Joshi, Sivan Toledo March 1998 ACM Transactions on Mathematical Software (TOMS), Volume 24 Issue 1

Publisher: ACM Press

Full text available: pdf(213.58 KB)


Additional Information: full citation, abstract, references, citings, index terms, review

This article describes the design, implementation, and evaluation of a parallel algorithm for the Cholesky factorization of symmetric banded matrices. The algorithm is part of IBM's parallel engineering and scientific subroutine library version 1.2 and is compatible with ScaLAPACK's banded solver. Analysis, as well as experiments on an IBM SP2 distributed-memory parallel computer, shows that the algorithm efficiently factors banded matrices with wide bandwidth. For example, a 31-mode SP2 fa ...

February 2005 ACM Transactions on Storage (TOS), Volume 1 Issue 1

Keywords: Banded matrices, Cholesky factorization, distributed memory, parallel algorithm

12 Consistent and automatic replica regeneration



Publisher: ACM Press

Full text available: pdf(372.24 KB) Additional Information: full citation, abstract, references, index terms

Reducing management costs and improving the availability of large-scale distributed systems require automatic replica regeneration, that is, creating new replicas in response to replica failures. A major challenge to regeneration is maintaining consistency when the replica group changes. Doing so is particularly difficult across the wide area where failure detection is complicated by network congestion and node overload. In this context, this article presents Om, the first read/write peer- ...

Keywords: Peer-to-peer storage systems, availability, consistency, regeneration, replication

13 Query evaluation techniques for large databases

Goetz Graefe June 1993 ACM Computing Surveys (CSUR), Volume 25 Issue 2

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(9.37 MB) terms, review

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

14 PocketLens: Toward a personal recommender system Bradley N. Miller, Joseph A. Konstan, John Riedl

July 2004 ACM Transactions on Information Systems (TOIS), Volume 22 Issue 3

Publisher: ACM Press

Full text available: pdf(1.10 MB) Additional Information: full citation, abstract, references, index terms

Recommender systems using collaborative filtering are a popular technique for reducing information overload and finding products to purchase. One limitation of current recommenders is that they are not portable. They can only run on large computers connected to the Internet. A second limitation is that they require the user to trust the owner of the recommender with personal preference data. Personal recommenders hold the promise of delivering high quality recommendations on palmtop computers, e ...

Keywords: Collaborative Filtering, Peer-to-Peer Networking, Privacy, Recommender



Systems

15 Analysis of local enumeration and storage schemes in HPF



Henk J. Sips, Kees van Reeuwijk, Will Denissen

January 1996 Proceedings of the 10th international conference on Supercomputing

Publisher: ACM Press

Full text available: pdf(909.88 KB) Additional Information: full citation, references, citings, index terms

16 Pursuing scalability for hypre's conceptual interfaces



Robert D. Falgout, Jim E. Jones, Ulrike Meier Yang

September 2005 ACM Transactions on Mathematical Software (TOMS), Volume 31 Issue 3

Publisher: ACM Press

Full text available: pdf(564.71 KB) Additional Information: full citation, abstract, references, index terms

The software library hypre provides high-performance preconditioners and solvers for the solution of large, sparse linear systems on massively parallel computers as well as conceptual interfaces that allow users to access the library in the way they naturally think about their problems. These interfaces include a stencil-based structured interface (Struct); a semistructured interface (semiStruct), which is appropriate for applications that are mostly structured, for example, block structu ...

Keywords: User interfaces, parallel programming, scalability

17 Decision Trees and Diagrams



Bernard M. E. Moret

December 1982 ACM Computing Surveys (CSUR), Volume 14 Issue 4

Publisher: ACM Press

Full text available: pdf(2.68 MB) Additional Information: full citation, references, citings, index terms

18 High performance Fortran language specification





CORPORATE Rice University

December 1993 ACM SIGPLAN Fortran Forum, Volume 12 Issue 4

Publisher: ACM Press

Additional Information: full citation, abstract, citings, index terms Full text available: pdf(5.69 MB)

(PART I)Fortran Forum is reprinting this High Performance Fortran Language Specification over several issues. The current issue is devoted to the first four chapters of the HPFF Language Specification. Remaining chapters of the HPFF Language Specification, and the HPFF Journal of Development, will be printed in installments in filture issues of Fortran Forum.

19 Proceedings of the SIGNUM conference on the programming environment for



development of numerical software

March 1979 ACM SIGNUM Newsletter, Volume 14 Issue 1

Publisher: ACM Press

Additional Information: full citation Full text available: pdf(5.02 MB)

20 Solution of large, dense symmetric generalized eigenvalue problems using



secondary storage

Roger G. Grimes, Horst D. Simon

September 1988 ACM Transactions on Mathematical Software (TOMS), Volume 14 Issue 3

Publisher: ACM Press

Full text available: pdf(1.12 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

This paper describes a new implementation of algorithms for solving large, dense symmetric eigen-problems AX = BX&Lgr;, where the matrices A and B are too large to fit in the central memory of the computer. Here A is assumed to be symmetric, and B symmetric positive definite. A combination of block Cholesky and block Householder transformations are used to reduce the problem to a symmetric ...

Results 1 - 20 of 200

Result page: **1** <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u>

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Mindows Media Player



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: • The ACM Digital Library • The Guide

join first table in a first storage and storing rows of second tab



THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used join first table in a first storage and storing rows of second table in a second storage

Found 112,882 of 167.655

Sort results by

Display

results

relevance

expanded form

Save results to a Binder

Search Tips

Open results in a new

Try an Advanced Search
Try this search in The ACM Guide

Results 41 - 60 of 200

window

Result page: previous 1 2 **3** 4 5 6 7 8 9 10

Relevance scale Relevance

Best 200 shown

41 Decision Decision City Alexandra in the Company

41 Revised Report of the Algorithmic Language Algol 68

A. van Wijngaarden

August 1981 ALGOL Bulletin, Issue Sup 47

Publisher: Computer History Museum

Full text available: pdf(9.20 MB)

Additional Information: full citation, index terms

42 Query processing: A characterization of the sensitivity of query optimization to





storage access cost parameters

Frederick R. Reiss, Tapas Kanungo

June 2003 Proceedings of the 2003 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(255.35 KB) Additional Information: full citation, abstract, references, index terms

Most relational query optimizers make use of information about the costs of accessing tuples and data structures on various storage devices. This information can at times be off by several orders of magnitude due to human error in configuration setup, sudden changes in load, or hardware failure. In this paper, we attempt to answer the following questions: • Are inaccurate access cost estimates likely to cause a typical query optimizer to choose a suboptimal query plan? • If an optimizer ...

Keywords: autonomic computing, computational geometry, databases, parametric query optimization, storage systems

43 The state of the art in distributed query processing



Donald Kossmann

December 2000 ACM Computing Surveys (CSUR), Volume 32 Issue 4

Publisher: ACM Press

Full text available: pdf(455.39 KB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> <u>terms</u>

Distributed data processing is becoming a reality. Businesses want to do it for many reasons, and they often must do it in order to stay competitive. While much of the infrastructure for distributed data processing is already there (e.g., modern network

technology), a number of issues make distributed data processing still a complex undertaking: (1) distributed systems can become very large, involving thousands of heterogeneous sites including PCs and mainframe server machines; (2) the stat ...

Keywords: caching, client-server databases, database application systems, dissemination-based information systems, economic models for query processing, middleware, multitier architectures, query execution, query optimization, replication, wrappers

44 XML data management and web discovery: Exploiting native XML indexing





techniques for XML retrieval in relational database systems

Felix Weigel, Klaus U. Schulz, Holger Meuss

November 2005 Proceedings of the 7th annual ACM international workshop on Web information and data management WIDM '05

Publisher: ACM Press

Full text available: pdf(569.93 KB) Additional Information: full citation, abstract, references, index terms

In XML retrieval, two distinct approaches have been established and pursued without much cross-fertilization taking place so far. On the one hand, native XML databases tailored to the semistructured data model have received considerable attention, and a wealth of index structures, join algorithms, tree encodings and query rewriting techniques for XML have been proposed. On the other hand, the question how to make XML fit the relational data model has been studied in great detail, giving r ...

Keywords: CADG, RCADG, RDBMS, RDBS, Relational CADG, XML indexing, XML retrieval, content-aware dataGuide, query evaluation, relational database, storage scheme

45 Research sessions: indexing and tuning: Transaction support for indexed summary





🜦 <u>views</u>

Goetz Graefe, Michael Zwilling

June 2004 Proceedings of the 2004 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(168.70 KB) Additional Information: full citation, abstract, references

Materialized views have become a standard technique for performance improvement in decision support databases and for a variety of monitoring purposes. In order to avoid inconsistencies and thus unpredictable query results, materialized views and their indexes should be maintained immediately within user transaction just like indexes on ordinary tables. Unfortunately, the smaller a materialized view is, the higher the concurrency contention between queries and updates as well as among concurrent ...

46 The STRIP rule system for efficiently maintaining derived data



Brad Adelberg, Hector Garcia-Molina, Jennifer Widom

June 1997 ACM SIGMOD Record, Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97, Volume 26 Issue 2

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.68 MB) <u>terms</u>

Derived data is maintained in a database system to correlate and summarize base data which records real world facts. As base data changes, derived data needs to be recomputed. This is often implemented by writing active rules that are triggered by changes to base data. In a system with rapidly changing base data, a database with a standard rule system may consume most of its resources running rules to recompute

data. This paper presents the rule system implemented as part of the STandard Re ...

47 Research session: query optimization and summarization: REED: robust, efficient filtering and event detection in sensor networks



Daniel J. Abadi, Samuel Madden, Wolfgang Lindner

August 2005 Proceedings of the 31st international conference on Very large data bases VLDB '05

Publisher: VLDB Endowment

Full text available: pdf(286.61 KB) Additional Information: full citation, abstract, references, index terms

This paper presents a set of algorithms for efficiently evaluating join queries over static data tables in sensor networks. We describe and evaluate three algorithms that take advantage of distributed join techniques. Our algorithms are capable of running in limited amounts of RAM, can distribute the storage burden over groups of nodes, and are tolerant to dropped packets and node failures. REED is thus suitable for a wide range of eventdetection applications that traditional sensor network dat ...

48 Research sessions: potpourri: Executing SQL over encrypted data in the database-



service-provider model

Hakan Hacigümüş, Bala İyer, Chen Li, Sharad Mehrotra

June 2002 Proceedings of the 2002 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(1.25 MB)

Additional Information: full citation, abstract, references, citings, index terms

Rapid advances in networking and Internet technologies have fueled the emergence of the "software as a service" model for enterprise computing. Successful examples of commercially viable software services include rent-a-spreadsheet, electronic mail services, general storage services, disaster protection services. "Database as a Service" model provides users power to create, store, modify, and retrieve data from anywhere in the world, as long as they have access to the Internet. It introduces sev ...

49 Caching: A self-managing data cache for edge-of-network web applications



Khalil Amiri, Sanghyun Park, Renu Tewari

November 2002 Proceedings of the eleventh international conference on Information and knowledge management

Publisher: ACM Press

Full text available: pdf(340.44 KB) Additional Information: full citation, abstract, references, index terms

Database caching at proxy servers enables dynamic content to be generated at the edge of the network, thereby improving the scalability and response time of web applications. The scale of deployment of edge servers coupled with the rising costs of their administration demand that such caching middleware be adaptive and self-managing. To achieve this, a cache must be dynamically populated and pruned based on the application query stream and access pattern. In this paper, we describe such a cache ...

Keywords: dynamic content, e-commerce, semantic caching

50 A Survey of Some Theoretical Aspects of Multiprocessing



J. L. Baer

January 1973 ACM Computing Surveys (CSUR), Volume 5 Issue 1

Publisher: ACM Press

Additional Information: full citation, references, citings, index terms Full text available: pdf(4.05 MB)

51 Efficient processing of spatial joins using R-trees



Thomas Brinkhoff, Hans-Peter Kriegel, Bernhard Seeger

June 1993 ACM SIGMOD Record, Proceedings of the 1993 ACM SIGMOD international conference on Management of data SIGMOD '93, Volume 22 Issue 2


Publisher: ACM Press

Full text available: pdf(1.48 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

Spatial joins are one of the most important operations for combining spatial objects of several relations. The efficient processing of a spatial join is extremely important since its execution time is superlinear in the number of spatial objects of the participating relations, and this number of objects may be very high. In this paper, we present a first detailed study of spatial join processing using R-trees, particularly R*-trees. R-trees are very suitable for supporting spatial queries a ...

52 Supporting valid-time indeterminacy



March 1998 ACM Transactions on Database Systems (TODS), Volume 23 Issue 1

Publisher: ACM Press

Full text available: pdf(516.09 KB)

Additional Information: full citation, abstract, references, citings, index

In valid-time indeterminacy it is known that an event stored in a database did in fact occur, but it is not known exactly when. In this paper we extend the SQL data model and query language to support valid-time indeterminacy. We represent the occurrence time of an event with a set of possible instants, delimiting when the event might have occurred, and a probability distribution over that set. We also describe query language constructs to retrieve informat ...

Keywords: SQL, TSQL2, incomplete information, indeterminacy, probabilistic information, temporal dtabase, valid-time database

⁵³ Query optimization in a memory-resident domain relational calculus database system





Kyu-Young Whang, Ravi Krishnamurthy

March 1990 ACM Transactions on Database Systems (TODS), Volume 15 Issue 1

Publisher: ACM Press

Full text available: pdf(2.46 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

We present techniques for optimizing queries in memory-resident database systems. Optimization techniques in memory-resident database systems differ significantly from those in conventional disk-resident database systems. In this paper we address the following aspects of query optimization in such systems and present specific solutions for them: (1) a new approach to developing a CPU-intensive cost model; (2) new optimization strategies for main-memory query processing; (3) new insight into ...

⁵⁴ A language and a physical organization technique for summary tables




Gultekin Ozsoyoglu, Z. Meral Ozsoyoglu, Francisco Mata May 1985 ACM SIGMOD Record, Proceedings of the 1985 ACM SIGMOD international conference on Management of data SIGMOD '85, Volume 14 Issue 4

Publisher: ACM Press

Full text available: 📆 pdf(1.20 MB)

Additional Information: full citation, references, citings, index terms

55 Multi-table joins through bitmapped join indices



September 1995 ACM SIGMOD Record, Volume 24 Issue 3

Publisher: ACM Press

Full text available: pdf(389.55 KB) Additional Information: full citation, abstract, citings, index terms

This technical note shows how to combine some well-known techniques to create a method that will efficiently execute common multi-table joins. We concentrate on a commonly occurring type of join known as a star-join, although the method presented will generalize to any type of multi-table join. A star-join consists of a central detail table with large cardinality, such as an orders table (where an order row contains a single purchase) with foreign keys that join to descriptive

56 Special issue on persistent object systems: Tigukat: a uniform behavioral objectbase management system



Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(2.78 MB) Additional Information: full citation, abstract, references, citings

We describe the TIGUKAT objectbase management system, which is under development at the Laboratory for Database Systems Research at the University of Alberta. TIGUKAT has a novel object model, whose identifying characteristics include a purely behavioral semantics and a uniform approach to objects. Everything in the system, including types, classes, collections, behaviors, and functions, as well as meta-information, is a first-class object with well-defined behavior. In this way, the model abstr ...

Keywords: database management, objectbase management, persistent storage system, reflective system

57 Query processing and optimization in Oracle Rdb

Gennady Antoshenkov, Mohamed Ziauddin

December 1996 The VLDB Journal — The International Journal on Very Large Data Bases. Volume 5 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(92.62 KB) Additional Information: full citation, abstract, citings, index terms

This paper contains an overview of the technology used in the query processing and optimization component of Oracle Rdb, a relational database management system originally developed by Digital Equipment Corporation and now under development by Oracle Corporation. Oracle Rdb is a production system that supports the most demanding database applications, runs on multiple platforms and in a variety of environments.

Keywords: Dynamic optimization, Optimizer, Query transformation, Relational database, Sampling

58 Practical minimal perfect hash functions for large databases

Edward A. Fox, Lenwood S. Heath, Qi Fan Chen, Amjad M. Daoud January 1992 Communications of the ACM, Volume 35 Issue 1

Publisher: ACM Press

Full text available: pdf(2.00 MB) Additional Information: full citation, references, citings, index terms

Keywords: CD-ROM, hashing, minimal perfect hash functions, perfect hash functions



⁵⁹ Ripple joins for online aggregation



Peter J. Haas, Joseph M. Hellerstein

June 1999 ACM SIGMOD Record, Proceedings of the 1999 ACM SIGMOD international conference on Management of data SIGMOD '99, Volume 28 Issue 2

Publisher: ACM Press

Full text available: pdf(1,78 MB)

Additional Information: full citation, abstract, references, citings, index terms

We present a new family of join algorithms, called ripple joins, for online processing of multi-table aggregation queries in a relational database management system (DBMS). Such queries arise naturally in interactive exploratory decision-support applications. Traditional offline join algorithms are designed to minimize the time to completion of the query. In contrast, ripple joins are designed to minimize the time until an acceptably precise estimate of the query result is availa ...

60 Sensor databases: The design of an acquisitional query processor for sensor





networks

Samuel Madden, Michael J. Franklin, Joseph M. Hellerstein, Wei Hong

June 2003 Proceedings of the 2003 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(485.52 KB)


Additional Information: full citation, abstract, references, citings, index


We discuss the design of an acquisitional query processor for data collection in sensor networks. Acquisitional issues are those that pertain to where, when, and how often data is physically acquired (sampled) and delivered to query processing operators. By focusing on the locations and costs of acquiring data, we are able to significantly reduce power consumption over traditional passive systems that assume the a priori existence of data. We discuss simple extensions to SQL for co ...

Results 41 - 60 of 200


Result page: <u>previous</u> <u>1</u> <u>2</u> **3** <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u> next

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2005 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us











Sign in



Froogle Local New! Groups News Images more » Advanced Search NCR and Teradata and distributing rows and d Preferences

The "AND" operator is unnecessary – we include all search terms by default. [details]

Web Results 41 - 50 of about 633 for NCR and Teradata and distributing rows and columns and tables. (1.

United States Patent Application: 0040236756

... system ("DBS") 100, such as a Teradata Active Data Warehousing System available from NCR Corporation. ... the storage of data and the distribution of table ... appft1.uspto.gov/.../ shu&RS=IN/shu - 38k - Supplemental Result - Cached - Similar pages

- Configuring and Tuning Databases on the Solaris™ Platform

... environment, specifically, a warehouse-based distribution operation. ... 6000k * scale factor rows. ... IBM DB2 and NCR Teradata, using similar technology, achieved ... wiki.nanthrax.net/.../OS/Solaris/configuring_and_tuning_databases_on_the_solaris_platform/ch27lev1sec4.html -13k - Supplemental Result - Cached - Similar pages

[PDF] ISO/IEC JTC 1/SC 32 N

File Format: PDF/Adobe Acrobat - View as HTML ... FYI REQUESTED ACTION DUE DATE Number of Pages 40 LANGUAGE USED English DISTRIBUTION P & L ... IBM, Informix, Microsoft, NCR/Teradata, Oracle, ... SQL tables onto XML ... www.jtc1sc32.org/sc32/jtc1sc32.nsf/Attachments/ D609402689C8841C88256AF8005F6CAA/\$FILE/32N0724-.PDF - Supplemental Result - Similar pages

IPDFI Microsoft PowerPoint - NCR-Partners2001-MetaIntegration.ppt

File Format: PDF/Adobe Acrobat - View as HTML Physical model (data types, table &. column definitions) directly from live. NCR Teradata databases via JDBC,. Complete Logical / Physical models ... www.metaintegration.net/Publications/ NCR-Partners2001-MetaIntegration.pdf - Similar pages

[PDF] SMM AWARDS SUPPLEMENT

File Format: PDF/Adobe Acrobat - View as HTML ... he made recommendations to all his distributors and retail ... Thomas Webster-Vice President of Marketing, Teradata, a division of NCR Corporation ... https://www.unisonindustries.com/ pdf/Featured%20Articles/SMMAWARDSSUPP.pdf - Supplemental Result -Similar pages

[PDF] Plan-Per-Tuple Optimization Solution - Parallel Execution of ... File Format: PDF/Adobe Acrobat - View as HTML exploring for the NCR Teradata Multimedia Database ... column itself relative to distribution of sizes of. videos. Sampling. implies running. a random ... www.vldb.org/conf/1998/p690.pdf - Similar pages

NCR: Teradata - Sequence number

I need to create a sequence number on a table. I was using row number logic ... That's why you have to order using column[s] with a good distribution. ... www.tek-tips.com/viewthread.cfm?qid=927608&page=1 - 22k - Cached - Similar pages

[PDF] Teradata Data Warehousing for Oracle Environments EB7651

File Format: PDF/Adobe Acrobat column and table levels. > Maintaining data using INSERT, ... Teradata and NCR are registered trademarks, and BYNET is a trademark of NCR Corporation. ... softlinebi.kiev.ua/data/upload/publication/ main/ru/143/migrationfromoracletotd.pdf - Similar pages [DOC] 1. Introduction

File Format: Microsoft Word - View as HTML

Over the last decade Teradata, Tandem, and a host of startup companies have ... common operator - it produces a row-and-column subset of a relational table. ...

research.microsoft.com/~gray/CacmParallelDB.doc - Similar pages

[More results from research.microsoft.com]

[PDF] beyond the data warehouse: the real-time revolution

File Format: PDF/Adobe Acrobat - View as HTML

The NCR Teradata response to a market moving toward requiring more ... technique,

table storage keys can be organized so that all the rows of a given entity ...

h71028.www7.hp.com/ERC/ downloads/BeyondtheDataWarehouse.pdf - Similar pages



Result Page: **Previous** 1 2 3 4 5 6 7 8 9 1011121314

NCR and Teradata and distributing r Search

Search within results | Language Tools | Search Tips

Google Home - Advertising Programs - Business Solutions - About Google

©2005 Google